


CERTIFICATE OF MAILING BY EXPRESS MAIL	
"EXPRESS MAIL" Mailing Label No <u>EL916516499US</u>	
Date of Deposit <u>November 14, 2001</u>	
I hereby certify that this paper or fee is being deposited with the U S Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1 10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Box Patent Application, Washington, D C 20231	
Type or Print Name <u>Marcy Overstreet</u>	
Signature	

SYSTEM AND METHOD FOR CONFIGURING OPTICAL CIRCUITS

Applicant(s): Andrea Fumagalli
James Cai

RELATED APPLICATION(S)

This application claims priority from and incorporates herein by reference the entire disclosure of U.S. Provisional Application Serial No. 60/248,265 filed November 14, 2000.

5 TECHNICAL FIELD

The present invention relates to the configuration of optical circuits, and more particularly, to the use of a multi-token control mechanism to configure optical circuits on demand.

BACKGROUND OF THE INVENTION

The insatiable appetite for Internet connectivity and network applications drives the current explosion of network traffic volume worldwide. It is expected that this exponential growth of traffic volume will continue in the foreseeable future. Optical fiber 5 communication technology based on Wavelength Division Multiplexing (WDM) has been employed as the major means to cope with the traffic volume growth. While WDM technology has already revolutionized the backbone network by enabling unprecedented increases in the leveraged capacity of a single fiber, a parallel paradigm shift is now taking place in the metropolitan network.

10 One of the most critical challenges in designing today's access and metropolitan networks is the fact that bandwidth demands have been consistently exceeding the most aggressive network planning predictions. In addition, the individual user's traffic burstiness makes static bandwidth reservation (e.g., SONET/SDH like) neither bandwidth efficient nor adequate to delay sensitive traffic. This situation has generated an increasing interest 15 towards all-optical networks that are capable of allocating network resources, i.e., bandwidth, in a dynamic way. Such networks must be able to reserve the necessary bandwidth on-demand to allow the transmission of a user's traffic burst. Once the burst transmission is completed, the reserved bandwidth is promptly released to be made available to other burst transmissions.

20 In order to be of practical use, the bandwidth on-demand concept requires few but fundamental features. Three of the features are:

- fast set-up time of the optical circuit (or lightpath);

- fair blocking probability irrespective of the lightpath span (or the number of fiber lines the lightpath is routed through)
- good bandwidth efficiency, i.e., the fraction of reserved bandwidth actually used to transmit data.

5

To understand how challenging it is to achieve these three features at once in the same architecture, one must observe that user requests for a lightpath are unpredictable and may occur simultaneously at distinct and geographically separated nodes. As a result, concurrent lightpath requests will compete to secure common resources, i.e., the available wavelengths 10 in the network. This may result in a number of reservation attempts being failed as they are blocked by other lightpath requests that book the resources first. In this scenario, it is thus possible to incur in long set-up times and unfair blocking probabilities that are a function of the lightpath span. Lightpaths with longer spans are more likely to be blocked since they require successful wavelength reservation on each and every fiber line they are routed 15 through.

Solutions so far proposed to solve the problem of routing and wavelength assignment (RWA) to establish lightpaths dynamically in a WDM ring can be categorized as centralized approaches and distributed approaches. With a centralized approach, the source node sends the request for a lightpath to a special node called controller. The controller keeps track of 20 the available network wavelengths and serves the node requests on a first-come-first-serve (FCFS) basis. The resource contention is resolved at the controller. On a unidirectional ring, latency of the signaling required between the source and the controller to set up and

eventually tear down the lightpath is proportional to the ring latency, i.e., round trip propagation time within the ring, and may considerably delay the set-up time and reduce bandwidth efficiency in metro applications.

With a distributed mechanism, every node solves the RWA problem for its own 5 newly requested lightpaths. One way to achieve this objective is to allow every node to keep track of network-wide wavelength availability. The RWA problem is solved based on shared global information. In another approach, each node makes use of a routing table for each wavelength which specifies the next hop and the cost associated with the shortest path to each destination on this wavelength. Since different nodes may concurrently try to assign the 10 same wavelength to distinct lightpath requests, both approaches require at least one round trip time from source to destination to be assured that their reservation was completed successfully. In a unidirectional ring this time equals the ring round trip time.

SUMMARY OF THE INVENTION

15 The present invention overcomes the foregoing and other problems with an optical network consisting of a source node and a destination node which are interconnected by a plurality wavelength, wherein each of the plurality of the wavelengths is associated with a particular channel. A token is associated with each of the plurality of wavelengths and indicates the availability of the wavelengths for supporting a lightpath. The source node is 20 configured to store a request for a lightpath between the source node and a destination node. Upon receipt of a token at the first node indicating an available space within the wavelength associated with the token, a request is selected from the queue using a best fit window

protocol. A connection is then established responsive to the selected request between the source node and the destination node.

The selection process would consist of determining whether any requests within the queue having expired soft deadlines and selecting a largest request having an expired soft 5 deadline which will fit within the available space of the wavelength for connection if any exist. If no soft deadline expirations are present, a largest request which will fit within the space available on the wavelength is selected. The selected request is used to establish a connection.

10 **BRIEF DESCRIPTION OF THE DRAWINGS**

A more complete understanding of the method and apparatus of the present invention may be obtained by reference to the following Detailed Description when taken in conjunction with the accompanying Drawings wherein:

15 FIGURE 1 generally illustrates the implementation of the system and method of the present invention;

FIGURE 2 illustrates source and destination nodes within an optical ring network;

FIGURE 3 illustrates the application of the best fit window approach of the present invention;

FIGURE 4 is a flow diagram illustrating the method of the present invention;

20 FIGURE 5 illustrates the achievable node throughput versus the average lightpath duration in a multiple ring latency for a centralized and distributed lightring analytical model;

FIGURE 6 illustrates the relationship between response time and throughput when using different sizes of a best fit window; and

FIGURE 7 illustrates achievable throughput for different burst sizes when using a different number of channels.

5

DETAILED DESCRIPTION

Referring now to the drawings, and more particularly to FIGURE 1, where there is provided a general illustration of the system of the present invention. The Lightring architecture of the present invention resorts to a unique distributed multi-token based control 10 wherein access to each wavelength 10 (channel) is controlled by a wavelength specific signaling-token 15 that is circulated among each node 20 on a ring 25 in a round robin fashion. For each data wavelength, a control message or token is continuously circulated among the nodes using the control channel. Tokens 15 regulate the access to the corresponding wavelength 10 and inform the source of the ring 25 available resources 15 (wavelengths). Tokens 15 bear resource availability information and broadcast this information to each node 20 in the ring network 25. This enables each node 20 of the ring network 25 to have an updated view of network resources. Upon reception of a token 15, a source node 20a with an outstanding lightpath request checks the available resources on the wavelength 10 associated with the token 15 and verifies if the outstanding lightpath can be 20 set up on that wavelength. If so, the token 15 is updated and passed onto the adjacent downstream node to inform all the other nodes 20 that a lightpath has been established and some resources have been reserved on that wavelength. A lightpath is set up between two

nodes 20 on a given wavelength only when a token 15 is acquired by a source node 20a. Similarly, for lightpath take-down, the token corresponding to the wavelength of the lightpath will be updated by the source to inform the other nodes of what resources have been freed. While circulating along the ring 25, tokens 15 broadcast lightpath status information 5 on other connections to each node 20 on the ring 25.

Referring now also to FIGURE 2, the network under consideration is a single fiber ring network 25 that connects N nodes 20. The network makes use of W data channels and one control channel, for a total of $W + 1$ wavelengths 10. Each wavelength supports one data channel. The optical signal on the control channel does not go through the node 20 and it is 10 separately handled by a control receiver 30 and a control transmitter 35. For each data channel a node 20 has one fixed control transmitter 30, one fixed control receiver 35 and one optical switch 40. This architecture allows the node 20 to transmit and receive message independently (and simultaneously) on any data channel. The on-off switches 45 within the optical switch 40 are used to control the flow of optical signals through the node 20 and 15 prevent signal re-circulation in the ring 25. A transmission buffer 55 is also provided at each node 20 to queue the generated packets prior to their transmission into the ring 25. The nodes 20 activities are regulated by an electronic controller 60 that determines the state for each on-off switch 45, the message transmission time, the wavelength used, and the reception 20 of the incoming messages. The electronic processing is done in parallel while the optical signal propagates through the fiber delay line 65 that connects the splitter 50 to a demultiplexer 70.

Once transmitted by the source node 20a, the message is removed from the network by the destination node 20b. Any uncollected section of the message, due to the setup of the optical switch at the destination node 20b, will make a round trip and be collected by the source node 20a. An optical copy of the message is obtained at every node 20 using a splitter 5 50, thus realizing a “broadcast and select” system. Only the intended destination(s) of the message actually receives the message.

Referring now back to FIGURE 1, in the present architecture, access to each wavelength 10 is controlled by a dedicated token 15 that is cyclically circulated among the nodes. Each data channel (i.e., wavelength) is associated with one token that is circulated 10 among the nodes 20 in the control channel and regulates the access to the corresponding channel. Thus, a total of W tokens 15 are available in the ring. A lightpath between a source node 20a and a destination node 20b may be set up and torn down only when a token 75 is acquired by the lightpath source node 20a. The token 15 is used to broadcast the wavelength status to all nodes 20 in a ring 25 and indicate whether there is available space on a 15 wavelength 10. Since only one node 20 at a time is allowed to make a reservation on each wavelength, the protocol of the present system achieves a “tell-and-go” reservation mechanism that is always successful. While circulating along the ring 25, tokens 15 broadcast the source 20a and destination node 20b of the newly established lightpath to all nodes 20 of the ring 25 so that no other node 20 will attempt to set up a lightpath on the same 20 wavelength 10 that overlaps in space with the one being established. Global lightpath status information of the moment is thus maintained on each node 20. This information can be stored in a memory (not shown) associated with each node 20.

If a strict first come first serve (FCFS) service policy is used on the message queue of each node 20, the system inclines to penalize the lightpath requests with longer spans when the offered load becomes high due to the space limit on the ring regarding a channel. Therefore, lightpath requests with long spans will hold up all the traffic behind it in the 5 queue. Based on this observation, a Best-Fit Window (BFW) mechanism is used to achieve better network throughput.

Contrary to all conventional wavelength assignment algorithms whereby an available (somehow optimal) wavelength is sought for each given lightpath request, the LightRing protocol seeks the lightpath request in the Beat-Fit-Window (BFW) of the transmission 10 queue that optimally fits the available space of the network on a given wavelength (identified by the arriving token) at the arrival time of the token. Any lightpath request in BFW of the transmission queue is a possible candidate to be transmitted based on the result of selection. The bandwidth efficiency achieved by the proposed reservation mechanism is proportional to the number of requests that the reservation mechanism can choose from, thus it is 15 proportional to the size of BFW. Complexity of the LightRing reservation mechanism is proportional to the size of BFW and not a function of the number of wavelengths (most of the existing reservation mechanisms have complexity that is proportional to the number of wavelengths). The LightRing reservation mechanism thus scales well when the number of wavelengths increase.

20 When a token 15 arrives, the best-fit message, which is the message with the longest span that can fit into the available space on the channel corresponding to the token, will be chosen and transmitted. Thus, as seen in FIGURE 3, where message requests 80-100 are

waiting in the queue 105, and a space 110 having a length N is available in a requested channel, message request 100 is selected for the space 110 because it is the request with the longest span capable of fitting in the available space.

In order to avoid starvation and guarantee fairness for the requests with different span lengths, a soft deadline is applied to each lightpath request so that the request will be dropped outside of BFW mechanism if the waiting time of the request is longer than a certain value. To maintain fairness of the system, requests are not dropped in the BFW mechanism. Instead, a FCFS protocol is applied to a request that has reached its soft deadline and the request is transmitted in the next available space. In other words, once a request gets into BFW, it has to be transmitted sooner or later.

Following is an explanation of each variable used in the protocol description.

r_i : i th lightpath request;

R_{TxQ} : $\{r_i | r_i \text{ in transmission buffer of a node}\}$;

R_{BFW} : $\{r_i | r_i \in R_{TxQ}, 0 \leq i < BFW\}$;

$R_{T \neq Q - BFW}$: $\{r_i | r_i \in R_{TxQ}, r_i \notin R_{BFW}\}$;

t : current time;

$t_i^{(a)}$: arrival time of r_i in R_{TxQ} (the time that r_i is inserted into the transmission queue);

$t_i^{(s)}$: beginning of the service time of r_i (the time r_i is removed from the

transmission queue);

$t_j^{(t)}$: arrival time for token j ;

$t_i^{(q)}$: time spent by r_i in $R_{T \neq Q}$;

$t_i^{(w)}$: time spent by r_i in R_{BFW} ;

$e^{(w)}(t)$: estimated average time spent by r_i in R_{BFW} as the function of t ;

d_j : soft deadline for serving lightpath request r_i ;

5 $d^{(w)}(t)$: soft deadline for r_i leaving R_{BFW} ; this value is based on the average time spent in the BFW .

R_{LATE} : $\{r_i | r_i \in R_{BFW}, t_i^{(w)} > d^{(w)}(t)\}$;

$l_j(i)$: the number of available hops left when r_i is placed into the space available on channel j ; negative number indicates the number of hops that r_i exceeds the 10 available space gap.

$Fit(R, \lambda_j)$: $\{r_i | r_i \in R, l_j(i) \geq 0\}$;

$BestFit(R, \lambda_j)$: $\{r_i | r_i \in R, l_j(k) \geq l_j(i), \text{for } \forall r_k \in R - r_i\}$;

$FCFS(R)$: $\{r_i | r_i \in R, t_j^{(a)}, \text{for } \forall r_k \in R - r_i\}$;

15 The following rules are applied to the protocol:

(1) only r_i in R_{BFW} can be served upon a token's arrival;

(2) only r_i in $R_{TxQ-BFW}$ may be dropped due to the soft deadline applied, so that once r_i gets into R_{BFW} , it has to be served eventually;

(3) arriving requests are dropped when no space is available in transmission 20 buffer;

The protocol used on a node can be described according to the following pseudocode.

However, it should be realized that other implementation of code are possible.

5 1. Upon arrival of token j , $t = t_j^{(t)}$ –

set up lightpath for request r_i

If $(Fit(R_{LATE}, \lambda_j) \neq \Phi)$ {

$r_i = FCFS(fit(R_{LATE}, \lambda_j))$

release token j with current lightpath info

10 transmit r_i

} else if $(Fit(R_{BFW}, \lambda_j) \neq \Phi)$ {

$r_i = BestFit(R_{BFW}, \lambda_j)$

release token j with current lightpath info

transmit r_i

15 } else

pass token j to next node;

2. Upon beginning of servicing of arrival of r_i , $t = t_i^{(j)}$ –

i) remove r_i from $R_{T \neq Q}$ and R_{BFW} ;

ii) set $e^{(w)}(t) = \beta e^{(w)} t_{i-1}^{(j)} + (1 - \beta) t_i^{(w)}$, where β is a system parameter of the estimator $e^{(w)}$, with value less than 1 but close to 1;

iii) set $d^{(w)}(t) = \alpha e^{(w)}(t)$, where $\alpha = \frac{d(w)(t)}{e(w)(t)}$ is the margin above average value for r_i spent in R_{BFW} · α is greater than 1 but close to 1;

5 iv) drop requests that past the soft deadline, which are $\{r_i \mid r_i \in R_{TQ-BFW}, t_i^{(q)} > d_{(q)} - e^{(w)}(t)\}$.

There are two fairness issues in the system of the present invention. The first issue is the fairness for the lightpath requests with different lengths in time. Since the incoming 10 lightpath request has exponential distribution on duration, this fairness issue is resolved automatically. This is true even when BFW is used because the request selection is totally independent from lightpath duration.

The second issue is the fairness for the lightpath requests with different spans (or 15 distances) on the ring. When FCFS policy is used, this fairness issue is also resolved automatically due to the uniform traffic distribution. But when BFW is used in the network, the simulation results confirm this fairness is no longer guaranteed. This is simply due to the fact that we do not randomly select requests in BFW regarding span length.

The way we tackle this problem is to apply a soft deadline $d^{(w)}(t)$ to each lightpath 20 request in BFW, where $d^{(w)(t)}$ is described above as a common filter with system parameter α and β . Any request with the time spent in BFW greater than $d^{(w)}(t)$ is considered late and

will be transmitted in the fashion of FCFS. The closer is the value of α to 1, the tighter is the constraint and in turn more observed is the fairness. The closer the value of β to 1, the slower the common filter reacts to input change.

If a hard deadline is applied to each lightpath request, the overdue request has to be 5 dropped no matter it is in BFW or not. Because of the non-random selection of request in BFW for transmission, if the overdue request is dropped in BFW, the blocking probability will no longer be the same for the requests with different lengths of span. Therefore the fairness can not be maintained any more. Based on this observation, we change to drop the overdue request right outside of BFW to maintain a soft deadline d_q . Since the estimated 10 average waiting time $e^{(w)}(t)$ is kept for record anyway, $d_q - e^{(w)}(t)$ can be used to check if we need to drop the request right outside of BFW when BFW has an empty spot to be filled. Therefore, in order to achieve the fairness o blocking probability only a soft deadline $d^{(q)}$ can be applied.

Referring now to FIGURE 4, there is illustrated a flow diagram generally describing 15 the process for assigning a request to a provided span within a channel. A token requesting establishment of a lightpath between a source node and a destination node is received at step 200. The request stored at step 205 within the nodes queue. An indication of an available span is received at step 210 from another token. Inquiry step 215 determines whether any 20 soft deadlines for any requests within a node queue have expired. If so, the first received request which will fit within the span and has an expired soft deadline is assigned to the span at step 220.

If inquiry step 215 determines that no request has exceeded their soft deadline, the span is compared to each message within the queue to select a best fit at 225. A best fit will comprise the request with the longest span which will fit within the available span length. The selected message is assigned to the span at step 230 so that a lightpath may be 5 established.

ANALYTICAL MODELS

In order to see the intrinsic difference between the normal centralized WDM ring with dynamic finite duration lightpath requests and the distributed architecture proposed 10 herein, a model for each case is presented. Due to the complexity of modeling LightRing with $BFW > 1$, we only consider the case when $BFW = 1$ and provide the simulation result for $BFW = 40$.

Our analysis extends the blocking probability model described in R. A. Barry and P. A. Humblet, "Models of Blocking Probability in All-Optical Networks with and without 15 Wavelength Changers," IEEE JSAC, Vol. 14, No. 5, June 1996, which is incorporated herein by reference, to capture the characteristics of dynamic traffic with finite duration. Barry's model introduces the qualitative behavior of the traffic for circuit-switched all-optical networks which can be used to calculate the blocking probability along a path. Yet the model does not cover the situation that lightpaths can be dynamically established and taken 20 down. The major variable Barry's model include P_i , the probability a lightpath ends and drops out at a node, and P_n , the probability a lightpath starts at a node on an available wavelength. The result is

$$P_n = \frac{\rho P_1}{1 - \rho(1 - P_1)} \quad (1)$$

where ρ is the utilization. The blocking probability without wavelength converter is

5

$$P_1 = \left| 1 - (1 - P_n)^R \right|^F \quad (2)$$

where H is the number of hope of the lightpath and F is the number wavelengths in each fiber.

10 In the case of unidirectional WDM ring, P_1 is I/N . In order to obtain the achievable throughput regardless the duration of the lightpaths, we can use the iteration technique due to the fact that $y = P_b^1 - P_b$, where P_b^1 is the result of iteration using P_b , is a monotonous function. Under maximized network load, here are the steps to find the blocking probability P_b in interaction working on $y - P_b$ plan:

15 1. set networth load $d = 1, P_{b0} = 0$ and $P_{b1} = 1$;

2. $\rho 0 = d(1 - P_{b0}); \rho 1 = d(1 - P_{b1})$;

3. get new P_{b0}^1 and P_{b1}^1 using Esq. 1 and 2;

4. $y0 = P_{b0}^1 - P_{b0}; y1 = P_{b1}^1 - P_{b1}$;

5. connect the point $(P_{b0}, y0)$ and $(P_{b1}, y1)$ with a straight line and find out the P_b that the line across the P_b axis;

6. $y = P_b^1 - P_{b1}$ where P_b^1 found using Eq. 1 and 2; if y has the same sign as $y0$,
 P_{b0} is replaced by P_b , otherwise P_{b1} is replaced by P_b ;
7. go back to step 2 until $|y|$ is less than a certain predetermined value;

5 Due to the establishment cost of lightpath with finite duration under both distributed
and centralized control mechanism, the real achievable throughput (thr) becomes

$$thr = E[\eta_a](1 - P_b) \quad (3)$$

10 where $E[\eta(\alpha)]$ is the average cost factor between the virtual throughput ($1 = P_b$) and the real
throughput (thr) for the average lightpath duration α .

CENTRALIZED APPROACH

With a centralized control mechanism, the source node sends the request for
15 a lightpath to a special node called the controller. The controller keeps track of the available
network wavelengths and serves the nodes' requests on a FCFS basis. Once the requested
lightpath is assigned a wavelength, the controller instructs the nodes that will set up the
optical add-drop multiplexers to establish the lightpath. The extra cost for setting up
lightpath is always one round trip delay. Therefore assuming the burst message length has
20 exponential distribution, the cost factor is

$$E[\eta(\alpha)] = \int_0^\infty \frac{t}{D+t} \frac{1}{a} e^{-t/a} dt \quad (4)$$

where D is the ring latency.

5 DISTRIBUTED APPROACH

In the proposed distributed Lightring protocol, the extra cost resides at the extra waiting time for the same token to come back to the source node when take down the lightpath. This is based on the assumption that no switching time is needed during the lightpath setup. Therefore the total time the lightpath in place is the multiple time of round trip delay that is immediately greater than the real lightpath duration α .

$$\begin{aligned}
 E[\mu(\alpha)] &= \sum_{n=0}^{\infty} \int_{nD}^{(n+1)D} \frac{t}{(n+1)D} \frac{1}{a} e^{-t/a} dt \\
 &= \sum_{n=0}^{\infty} \frac{a}{D(n+1)} (e^{-nD/a} - e^{-(n+1)D/a}) \\
 &+ \sum_{n=0}^{\infty} \frac{1}{n+1} (ne^{-nD/a} - (n+1)e^{-(n+1)D/a}) \\
 &= 1 + (e^{D/a} + \frac{a(1-e^{D/a})}{D}) \ln(1 - e^{-D/a}) \quad (5)
 \end{aligned}$$

15

18

PERFORMANCE RESULTS

The performance results presented are produced from the simulation model implemented in C++ and the analytical model described above. Unless indicated explicitly, 5 the network under consideration is a WDM ring with 32 wavelengths and 16 nodes evenly distributed over 80 km of fiber. Each wavelength supports a fixed transmission rate of 10 Gbps. For demonstration purposes, we assume the network traffic has Poisson arrival rate and lightpath duration is exponentially distributed. Traffic is uniformly distributed, meaning that the source and the destination nodes of a newly generated message are randomly chosen.

10 FIGURE 5 depicts the achievable node throughput vs. the average lightpath duration in the multiple of the ring latency for the centralized and distributed LightRing analytical model presented above and the simulation result of the LightRing protocol with BFW = 1 and BFW = 40. The distributed model has BFW = 1, it fairly closely matches the simulation result with BFW = 1. As we can see when burst is not too large, LightRing clearly 15 outperforms the centralized approach. Theoretically, the two curves will converge when the burst size approaches infinity. Also when BFW size increases, bandwidth efficiency is also improved.

FIGURE 6 shows the relationship between the response time and throughput when using different sizes of BFW. Response time is defined as the summation of the waiting time 20 in queue and the transmission time. The average message length is 10 Mbit. The performance improvement of using larger BFW occurs under medium to heavy load. The improvement is the most obvious when BFW first picks up and becomes less obvious later.

As FIGURE 6 shows, response time and control complexity can be traded for bandwidth efficiency by varying the BFW size. It is also noticed that when the network load is not too heavy, the response time can be well below the summation of the ring latency and the average burst duration (in this case, it is 1.4 msec). In other words, the time to establish a 5 lightpath can be well below the ring latency as opposed to the case for existing centralized and distributed reservation mechanisms that needs at least the ring round trip time.

Based on the LightRing protocol, FIGURE 7 shows the achievable throughput for different burst sizes when using a different number of channels. The result is based on the analytical model described above. The total bandwidth is fixed to 80 Gbps, so that when the 10 number of channels increases, the transmission rate for each individual channel will decrease. That implies lower costs for the transmitter and receiver, assuming cost has more than linear growth while transmission speed increases. But more importantly FIGURE 7 indicates better bandwidth efficiency when the number of channels increase. This is due to the fact that a node acquires tokens more frequently and more space to set up a lightpath.

15 Finally, the blocking probability for the lightpaths with a different number of hops is completely fair in LightRing due to the uniform traffic, and the fact that late messages are removed from the transmission queue only outside of BFW.

The LightRing architecture was presented in which a multi-token based reservation mechanism is used to set up lightpaths on-demand. By performing a tell-and-go reservation 20 of the wavelengths, the LightRing approach yields fast set-up time and efficient bandwidth utilization even in presence of relatively short bursts of data, e.g., bursts whose transmission time is 1ms in a 80 km ring.

Among other interesting features, the LightRing architecture is compatible with optical packet switching, and its performance improves with the number of wavelengths, consistently with the current trend of optical technologies. Complexity of the reservation mechanism is not a function of the number of wavelengths, and can be varied to trade 5 response time for bandwidth efficiency. Finally, the LightRing approach is compatible with emerging protocols for bandwidth reservation in the optical layer, e.g., MP λ S, and yields fair blocking probability irrespective of the lightpath span.

The previous description is of a preferred embodiment for implementing the invention, and the scope of the invention should not necessarily be limited by this 10 description. The scope of the present invention is instead defined by the following claims.